arXiv:1606.03196v1 [cs.IT] 10 Jun 2016

Phase Retrieval via Incremental Truncated Wirtinger Flow

Ritesh Kolte
Department of Electrical Engineering
Stanford University
Stanford, CA 94305 USA
ritesh.kolte@gmail.com

Ayfer Özgür
Department of Electrical Engineering
Stanford University
Stanford, CA 94305 USA
aozgur@stanford.edu

Abstract

In the phase retrieval problem, an unknown vector is to be recovered given quadratic measurements. This problem has received considerable attention in recent times. In this paper, we present an algorithm to solve a nonconvex formulation of the phase retrieval problem, that we call *Incremental Truncated Wirtinger Flow*. Given random Gaussian sensing vectors, we prove that it converges linearly to the solution, with an optimal sample complexity. We also provide stability guarantees of the algorithm under noisy measurements. Performance and comparisons with existing algorithms are illustrated via numerical experiments on simulated and real data, with both random and structured sensing vectors.

1 Introduction

Let $\boldsymbol{x} \in \mathbb{C}^n$ be an arbitrary unknown vector. We consider the problem of recovering $\boldsymbol{x}$, given data which is m quadratic measurements

$$y_i = |\boldsymbol{a}_i^* \boldsymbol{x}|^2, \quad 1 \leq i \leq m, \tag{1}$$

where the sensing vectors $\boldsymbol{a}_i \in \mathbb{C}^n$ are known. The problem of solving this quadratic set of equations is quite general: special cases include the phase retrieval problem which arises in many imaging applications due to physical limitations of sensors, as well as the NP-complete stone problem in combinatorial optimization.

Despite the hardness, impressive theoretical and empirical performance guarantees have been obtained recently by making some assumptions about the model, such as assuming the sensing vectors to be i.i.d. samples from, say $\mathcal{N}(0, I)$ or $\mathcal{CN}(0, I)$. The first work in this direction was [1], which employed the squared loss function and attempted to solve

$$\underset{\boldsymbol{z} \in \mathbb{R}^n}{\text{minimize}} \frac{1}{m} \sum_{i=1}^{m} \ell(y_i, |\boldsymbol{a}_i^* \boldsymbol{z}|^2), \tag{2}$$

where $\ell(y_i, |\boldsymbol{a}_i^* \boldsymbol{z}|^2) = (y_i - |\boldsymbol{a}_i^* \boldsymbol{z}|^2)^2$. The non-convexity of this problem was addressed by performing a lifting step (expressing the problem in terms of $\boldsymbol{z}\boldsymbol{z}^T$) followed by a convex relaxation. This algorithm recovers $\boldsymbol{x}$ with high probability if $m = O(n)$, [2]. However, the memory requirements and computational complexity make it prohibitive for problems of large dimension.

Follow-up works such as [3], [4], [5] made progress on the computational complexity front. All of these works developed algorithms to solve the problem (2) directly without performing the lifting step, using a first-order optimization method for iterative refinement, after an appropriate initialization. The work [5] also provided stability guarantees if the measurements are corrupted with noise.

However, each iteration in these algorithms requires one pass through the entire data. This can be highly undesirable when the dimensions of the problem are large, since a single update can require a

large amount of time, in part due to the communication delays introduced when the entire data does not fit in the available memory. Large dimensions naturally arise in the phase retrieval problem since the object of interest usually represents an image, so n is the product of the image dimensions.

In this paper, we build on the idea of Truncated Wirtinger Flow (TWF) from [5] and modify it to obtain the Incremental Truncated Wirtinger Flow (ITWF). By *incremental*, we mean that each iteration of the algorithm only accesses one randomly chosen data point, i.e. one sensing vector and the corresponding measurement. Thus, each iteration of ITWF is cheaper than that of TWF by a factor m, similar to what happens, e.g., by going from full gradient descent to stochastic gradient descent (SGD) in the case of standard empirical risk minimization problems. Unfortunately, this benefit is not obtained readily for the problem at hand, since the truncation performed at each iteration in [5] makes use of a threshold that is a function of all the sensing vectors and measurements. Excluding the truncation leads to a severe hit in the performance as observed in [5]. Furthermore, similar to SGD, sampling one data point instead of all data points introduces variance in the descent direction. The main contribution of this paper is the design of the incremental method ITWF that matches the excellent performance of TWF in terms of the statistical complexity, computational complexity and robustness to noisy measurements despite being an incremental method. In fact, our numerical experiments demonstrate that ITWF far surpasses TWF on the computational complexity front.

Algorithm	Sample complexity m	Computational complexity (stage II)
AltMinPhase [4]	$O(n \log^3 n)$	$O(n^2 \log^3(n/\epsilon))$
Wirtinger Flow [3]	$O(n \log n)$	$O\left(n^3 \log(n/\epsilon)\right)$
Truncated Wirtinger Flow [5]	$O(n)$	$O\left(n^2 \log(1/\epsilon)\right)$
This paper	$O(n)$	$O\left(n^2 \log(1/\epsilon)\right)$

Table 1: Performance of Algorithms

Table 1 provides a comparison of algorithms. Though our original intention was to simply develop an incremental version of TWF to allow efficient handling of large data, we find that ITWF provides a remarkable speedup as compared to TWF, even at $n \approx 1000$. To provide the reader an idea about the speed-up, we mention here some observations from the numerical experiments in Section 4. In Example 2, after an initialization stage requiring 10 passes through the data, the second stage of TWF requires at least 120 further passes to get a high accuracy solution. In contrast, the second stage of ITWF recovers a high accuracy solution in less than 15 passes. Example 3 presents an even more compelling case. Here, after 50 passes through the data for initialization, the solution returned by TWF after 50 further passes can instead be obtained by making 3 passes using ITWF.

1.1 Related Work

There has also been a lot of recent work on different formulations of the phase retrieval problem. In the sparse phase retrieval problem, it is additionally assumed that the vector x has only a few non-zero entries, and algorithms are sought for which the sample complexity and computational complexity have optimal dependence on the number of non-zero entries. The interested reader is referred to works such as [4], [6], [7]. Phase retrieval under the assumption of structured sensing vectors has also been a topic of interest, e.g. coded diffraction patterns [8], STFT [9]. We also present numerical experiments in this paper demonstrating the performance of ITWF when the sensing vectors are structured. Recent developments on different tractable formulations of the phase retrieval problem have been compiled in the survey article [10]. Finally, while we focus on devising an incremental update for the iterative refinement stage of TWF, the initialization stage can also be made incremental in a straightforward manner by using incremental algorithms for PCA, such as the algorithm from [11].

2 Main Idea

In the following section, we first describe the algorithm TWF from [5], and then describe the proposed algorithm ITWF.

Since we can only hope to recover the solution upto a global phase, we define $\text{dist}(\boldsymbol{z}, \boldsymbol{x})$ to be $\min_{\varphi \in [0,2\pi)} \|e^{-j\varphi}\boldsymbol{z} - \boldsymbol{x}\|$, and we implicitly assume in the remainder of the paper that $\boldsymbol{z}$ is $e^{-j\varphi(\boldsymbol{z})}\boldsymbol{z}$,

where $\varphi(\boldsymbol{z})$ is the argument of the above minimization problem. Then, we use $\boldsymbol{h}$ to denote $\boldsymbol{z} - \boldsymbol{x}$. Thus, $\|\boldsymbol{h}\| = \text{dist}(\boldsymbol{z}, \boldsymbol{x})$.

2.1 Truncated Wirtinger Flow

The loss function used in [5] is $\ell(y, \hat{y}) = \hat{y} - y \log \hat{y}$, which would correspond to maximizing the log likelihood if the measurements were assumed to arise from a Poisson noise model. The TWF algorithm has two stages described below, and requires choosing the constants α_y, μ, α_z^{lb}, α_z^{ub}, α_h.

(1) Truncated Spectral Initialization
Set $\boldsymbol{z}^{(0)}$ to be the principal eigenvector (appropriately scaled) of

$$\frac{1}{m}\sum_{i=1}^{m} y_i \boldsymbol{a}_i \boldsymbol{a}_i^* \mathbb{1}_{\left\{y_i \le \alpha_y^2 \left(\frac{1}{m}\sum_{i=1}^{m} y_i\right)\right\}},$$

with α_y set to, say 3. The point $\boldsymbol{z}^{(0)}$ satisfies

$$\text{dist}(\boldsymbol{z}^{(0)}, \boldsymbol{x}) \le \delta \|\boldsymbol{x}\|, \tag{3}$$

for any $\delta > 0$, as long as m/n exceeds a sufficiently large constant. For intuition and proof of this fact, we refer the reader to [5].

(2) Truncated Wirtinger Flow
For $t = 0, 1, \ldots, T-1$, perform the update

$$\boldsymbol{z}^{(t+1)} = \boldsymbol{z}^{(t)} - \frac{\mu}{m}\sum_{i=1}^{m} \nabla \ell(y_i, |\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2) \cdot \mathbb{1}_{\mathcal{E}_{1,t}^i \cap \mathcal{E}_{2,t}^i}, \tag{4}$$

where $\nabla \ell(y_i, |\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2)$ denotes the Wirtinger derivative of with respect to $\boldsymbol{z}$:

$$\nabla \ell(y_i, |\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2) = \frac{2|\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2 - 2y_i}{\boldsymbol{z}^{(t)*} \boldsymbol{a}_i} \boldsymbol{a}_i,$$

and the events $\mathcal{E}_{1,t}^i$ and $\mathcal{E}_{2,t}^i$ are defined as follows,

$$\mathcal{E}_{1,t}^i := \left\{ \alpha_z^{\text{lb}} \le \frac{|\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|}{\|\boldsymbol{z}^{(t)}\|} \le \alpha_z^{\text{ub}} \right\},$$

$$\mathcal{E}_{2,t}^i := \left\{ \left| y_i - |\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2 \right| \le \frac{\alpha_h}{m} \sum_{i=1}^{m} \left| y_i - |\boldsymbol{a}_i^* \boldsymbol{z}^{(t)}|^2 \right| \right\}. \tag{5}$$

The idea is to perform an update similar to gradient descent. The truncation results in dropping those indices for which the numerator and denominator magnitudes are very different from their expected values respectively, since such terms can exert a large atypical influence causing the full gradient to point in an undesirable direction.

2.2 Incremental Truncated Wirtinger Flow

The straightforward way of turning the above algorithm into an incremental one would be to replace (4) by:

$$\boldsymbol{z}^{(t+1)} = \boldsymbol{z}^{(t)} - \mu \nabla \ell(y_{i_t}, |\boldsymbol{a}_{i_t}^* \boldsymbol{z}^{(t)}|^2) \cdot \mathbb{1}_{\mathcal{E}_{1,t}^{i_t} \cap \mathcal{E}_{2,t}^{i_t}}, \tag{6}$$

where i_t would be chosen uniformly at random from $\{1, 2, \ldots, m\}$. However, as can be seen from (5), checking if $\mathcal{E}_{2,t}^i$ has occurred for any i requires a full pass through the entire data. Thus, (6) is as costly as (4). To address this difficulty, we replace $\mathcal{E}_{2,t}^i$ by the event $\mathcal{E}_3^i$, which is defined as

$$\mathcal{E}_3^i := \left\{ \frac{y_i}{\frac{1}{m}\sum_{i=1}^{m} y_i} \le (\alpha_x)^2 \right\},$$

where α_x is a constant to be chosen. This is amenable for an incremental update, which we call Incremental Truncated Wirtinger Flow (ITWF):

$$\boldsymbol{z}^{(t+1)} = \boldsymbol{z}^{(t)} - \mu \nabla \ell(y_{i_t}, |\boldsymbol{a}_{i_t}^* \boldsymbol{z}^{(t)}|^2) \cdot \mathbb{1}_{\mathcal{E}_{1,t}^{i_t} \cap \mathcal{E}_3^{i_t}}. \tag{7}$$

Data: Measurements and sampling vectors $\{y_i, \boldsymbol{a}_i\}_{1\le i\le m}$
Result: $\boldsymbol{z}^{(T)}$
Stage I (initialization):
Set $\boldsymbol{z}^{(0)}$ to be $\sqrt{\frac{1}{m}\sum_{i=1}^m y_i}\cdot \boldsymbol{z}_{\text{init}}$, where $\boldsymbol{z}_{\text{init}}$ is the principal eigenvector of $\frac{1}{m}\sum_{i=1}^m y_i \boldsymbol{a}_i\boldsymbol{a}_i^* \mathbb{1}_{\left\{y_i\le \alpha_y^2\left(\frac{1}{m}\sum_{i=1}^m y_i\right)\right\}}$.
Stage II:
for $t = 1, 2, \ldots, T-1$ **do**
  Sample i_t uniformly at random from $\{1, 2, \ldots, m\}$

$$\boldsymbol{z}^{(t+1)} = \boldsymbol{z}^{(t)} - \mu\nabla \ell(y_{i_t}, |\boldsymbol{a}_{i_t}^*\boldsymbol{z}^{(t)}|^2)\cdot \mathbb{1}_{\mathcal{E}_{1,t}^{i_t}\cap \mathcal{E}_3^{i_t}}$$

end

Algorithm 1: Incremental Truncated Wirtinger Flow

Note in fact that not only does $\mathcal{E}_3^i$ not depend on the entire data, but it is non-adaptive, i.e. it does not depend on $\boldsymbol{z}^{(t)}$. So, $\mathcal{E}_3^i$ simply discards measurements in which $|\boldsymbol{a}_i^T\boldsymbol{x}|$ deviates a lot from $\|\boldsymbol{x}\|$, similar to the rule used during the initialization stage.

The main implication of (5) used in [5] is that the magnitude of the term $\boldsymbol{a}_i^T\boldsymbol{h}$ can be controlled, which turns out to be crucial in obtaining the optimal sample complexity $m = O(n)$. Since ITWF cannot employ this truncation rule, we are not able to control the magnitude of $\boldsymbol{a}_i^T\boldsymbol{h}$. However, replacing (5) by $\mathcal{E}_3^i$ does not result in any deterioration in the sample complexity and computational complexity.

Furthermore, we would like to point out an empirical observation that we did not observe a significant difference in the numerical experiments even if we only employed truncation based on $\mathcal{E}_{1,t}^i$, i.e. by excluding $\mathcal{E}_3^i$! In fact, we also observed that TWF also performs similarly whether $\mathcal{E}_{2,t}^i$ is included or not. This suggests that while the theoretical analysis of TWF and ITWF require these additional truncation events, it might be possible to prove the convergence results via a different line of analysis that does not require introducing these events. However, we observed that truncation based on $\mathcal{E}_{1,t}^i$ is indeed crucial in the numerical experiments.

3 Main Results and Discussion

We focus on the real valued case for simplicity of exposition, and for concreteness, we fix the constants α_z^{lb}, α_z^{ub} and α_x to be 0.3, 5 and 5 respectively. The following two theorems are counterparts of the two main results in [5]. The first result, Theorem 1, focuses on the noiseless case (1) and provides a linear convergence guarantee.

Theorem 1. *Under noiseless measurements* (1) *with* $\{\boldsymbol{a}_i\}_{i=1}^m \sim \mathcal{N}(0, I)$ *independent, there exist universal constants* $C, c_0, c_1, c_2 > 0$ *and* $0 < \rho, \nu < 1$*, such that with probability at least* $1 - Cm\exp(-c_1 n)$ *and* $\mu = c_2/n$*, the iterates in Algorithm 1 satisfy*

$$\mathbb{E}_{\mathcal{I}^t}\left[\text{dist}^2(\boldsymbol{z}^{(t)}, \boldsymbol{x})\right] \le \nu\left(1 - \frac{\rho}{n}\right)^t \|\boldsymbol{x}\|^2, \tag{8}$$

if $m \ge c_0 n$.

Thus, the (mean squared error) MSE is reduced by a factor $(1-\rho/n)^m$ after one pass through the data.

In the above statement, $\mathbb{E}_{\mathcal{I}^t}[\cdot]$ denotes the expectation with respect to algorithm randomness $\mathcal{I}^t = \{i_1, i_2, \ldots, i_{t-1}\}$. More formally, it denotes the expectation conditioned on the data randomness $\{\boldsymbol{a}_i\}_{i=1}^m$. To keep expressions simple, we will follow the convention that $\mathbb{E}_X[\cdot]$ denotes expectation conditioned on all random variables except X.

Note that the expectation is only with respect to the algorithm randomness, not with respect to the data randomness $\{\boldsymbol{a}_i\}_{i=1}^m$. This is crucial since the data could be provided as it is, thus necessitating the need for a convergence guarantee that holds with high probability with respect to the data randomness.

Since the convergence is linear, a simple application of Markov's inequality already provides a strong convergence guarantee in which the high probability also refers to the algorithm randomness.

Also note that linear convergence is achieved for our setup even with an incremental method since the effect of the variance of the stochastic gradient can be controlled without having to choose a step-size that decreases with iteration, while in general, incremental methods suffer from slow convergence unless variance-reducing modifications [12] are used. The reason for this is explained in the next subsection.

Instead of (1), if the measurements are noisy such that

$$y_i = |\boldsymbol{a}_i^* \boldsymbol{x}|^2 + \eta_i, \quad 1 \le i \le m, \tag{9}$$

where η_i denotes the noise term (need not be stochastic), then we have the following theorem which provides a stability guarantee.

Theorem 2. *Under noisy measurements* (9) *with* $\{\boldsymbol{a}_i\}_{i=1}^m \sim \mathcal{N}(0, I)$ *independent and the noise satisfying* $\|\boldsymbol{\eta}\|_\infty \le \epsilon_\eta \|\boldsymbol{x}\|^2$ *for some small constant* $\epsilon_\eta > 0$*, there exist universal constants* $C, c_0, c_1, c_2 > 0$ *and* $0 < \rho, \nu < 1$*, such that with probability at least* $1 - Cm\exp(-c_1 n)$ *and* $\mu = c_2/n$*, the iterates in Algorithm 1 satisfy*

$$\mathbb{E}_{\mathcal{I}^t}\left[\text{dist}^2(\boldsymbol{z}^{(t)}, \boldsymbol{x})\right] \lesssim \frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{x}\|^2} + \left(1 - \frac{\rho}{n}\right)^t \|\boldsymbol{x}\|^2, \tag{10}$$

if $m \ge c_0 n$.

As described in [5], this result can be applied to the case when the measurements are obtained independently according to a Poisson noise model:

$$y_i \sim \text{Poisson}(|\boldsymbol{a}_i^* \boldsymbol{x}|^2) \quad 1 \le i \le m, \tag{11}$$

and the solution satisfies $\|\boldsymbol{x}\|^2 \ge \log^3 m$ (required to ensure that the condition $\|\boldsymbol{\eta}\|_\infty \le \epsilon_\eta \|\boldsymbol{x}\|^2$ holds with high probability).

3.1 Intuition

Consider the noiseless case. The reason why we can expect linear convergence from ITWF can be understood by the following thought experiment. Say, after the t^{th} iteration of ITWF, we had the luxury of obtaining a new measurement by sampling $\boldsymbol{a}_t$ independently from the population distribution $\mathcal{N}(0, I)$, instead of being restricted to sample from the empirical distribution $\{\boldsymbol{a}_i\}_{i=1}^m$.

Let $\boldsymbol{h}$ denote $\boldsymbol{z}^{(t)} - \boldsymbol{x}$, and $\mathcal{E}_t = \mathcal{E}_{1,t} \cap \mathcal{E}_3 = \left\{\alpha_z^{\text{lb}} \le \frac{|\boldsymbol{a}_t^* \boldsymbol{z}^{(t)}|}{\|\boldsymbol{z}^{(t)}\|} \le \alpha_z^{\text{ub}}\right\} \cap \left\{\frac{|\boldsymbol{a}_t^T \boldsymbol{x}|}{\|\boldsymbol{x}\|} \le \alpha_x\right\}$ denote the overall truncation event. Then the expected distance to the optimal solution after performing the ITWF update (7) using the new independent measurement, is

$$\begin{aligned}\mathbb{E}_t\left[\text{dist}^2(\boldsymbol{z}^{(t+1)}, \boldsymbol{x})\right] = \|\boldsymbol{h}\|^2 - 4\mu\mathbb{E}_t\left[|\boldsymbol{a}_t^T \boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_t^T \boldsymbol{x}}{\boldsymbol{a}_t^T \boldsymbol{z}^{(t)}}\right) \mathbb{1}_{\mathcal{E}_t}\right] \\ + 4\mu^2 \mathbb{E}_t\left[\|\boldsymbol{a}_t\|^2 |\boldsymbol{a}_t^T \boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_t^T \boldsymbol{x}}{\boldsymbol{a}_t^T \boldsymbol{z}^{(t)}}\right)^2 \mathbb{1}_{\mathcal{E}_t}\right]\end{aligned} \tag{12}$$

Since $\boldsymbol{a}_t \sim \mathcal{N}(0, I)$ independent of $\boldsymbol{h}, \boldsymbol{z}, \boldsymbol{x}$, it is not difficult to show that if $\|\boldsymbol{h}\|$ is sufficiently small compared to $\|\boldsymbol{x}\|$,

$$\mathbb{E}_t\left[|\boldsymbol{a}_t^T \boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_t^T \boldsymbol{x}}{\boldsymbol{a}_t^T \boldsymbol{z}}\right) \mathbb{1}_{\mathcal{E}_t}\right] \ge (1 - c_1)\|\boldsymbol{h}\|^2, \ \mathbb{E}_t\left[\|\boldsymbol{a}_t\|^2 |\boldsymbol{a}_t^T \boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_t^T \boldsymbol{x}}{\boldsymbol{a}_t^T \boldsymbol{z}}\right)^2 \mathbb{1}_{\mathcal{E}_t}\right] \lesssim n\|\boldsymbol{h}\|^2,$$

for an appropriate constant c_1. As can be observed from the latter inequality, the variance of the stochastic gradient is proportional to $\|\boldsymbol{h}\|^2$. As a result, the variance is not only bounded, but reduces as we get closer to the solution, which allows us to choose a non-diminishing step-size $\mu = \Theta\left(\frac{1}{n}\right)$, resulting in linear convergence:

$$\mathbb{E}_t\left[\text{dist}^2(\boldsymbol{z}^{(t+1)}, \boldsymbol{x})\right] \le \left(1 - \frac{\rho}{n}\right) \text{dist}^2(\boldsymbol{z}^{(t)}, \boldsymbol{x}).$$

In the above thought experiment, we had the luxury of an infinite amount of measurements at our disposal. Back in reality, we have the restriction of using a finite set of m measurements. Since we need to keep sampling from this set to simulate the stochastic gradient descent from our thought experiment, the proof of Theorem 1 mainly involves showing that we can draw similar conclusions despite the finiteness of the data set and the resulting dependence between the sensing vectors and the path traversed by the algorithm.

4 Numerical Experiments

To provide evidence for the performance of ITWF and comparisons to existing algorithms, we provide numerous examples in this section.

4.1 Example 1: Sample complexity

We compare the sample complexity of TWF and ITWF empirically when the sensing vectors $\boldsymbol{a}_i$ are i.i.d. $\mathcal{N}(0, I)$. The dimension n is chosen to be 1000, and the number of measurements m is varied from $2n$ to $6n$, and success is declared if the relative root mean squared error $\frac{\text{dist}(\boldsymbol{z},\boldsymbol{x})}{\|\boldsymbol{x}\|}$ is less than 10^{-5} within 1000 passes through the data. The initialization uses 50 truncated power iterations. It can be seen from Figure 1, which is obtained by averaging over 100 Monte Carlo trials at each value of m, that the empirical success rate of ITWF is comparable with that of TWF, even slightly better.

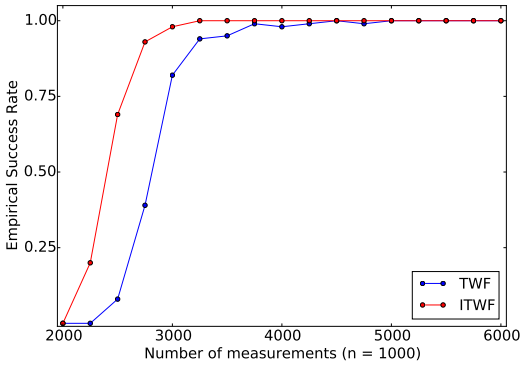

Figure 1: Empirical Success Rate

4.2 Example 2: Computational Complexity (Stage II)

We consider the same setup as the previous example, with m fixed to be 8000. Figure 2 shows the relative root mean squared error $\frac{\text{dist}(\boldsymbol{z},\boldsymbol{x})}{\|\boldsymbol{x}\|}$ of three algorithms, each run with the best step size, as a function of the number of passes through the data. All algorithms were initialized using 10 power iterations. The black line corresponds to the without-replacement variant of ITWF, in which all data points are visited exactly once in every block of m iterations, each time in an independently chosen random order. The performance of the without-replacement variant is marginally better than the with-replacement variant, as has also been observed in numerous other contexts. However, the main point of the example is to show that ITWF offers substantial improvements in computational complexity over TWF. Owing to the better performance of the without-replacement variant, we adopt this sampling method for the remainder of the numerical experiments.

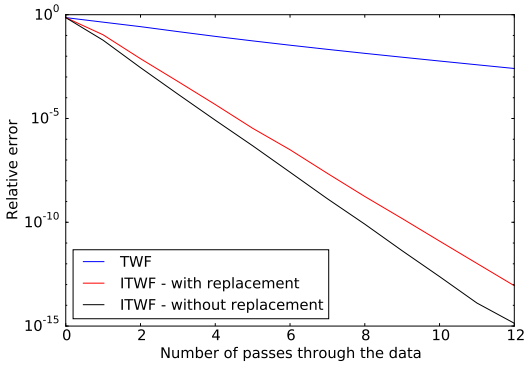

Figure 2: Relative Error

4.3 Example 3: Structured Sensing Vectors

To demonstrate the performance of the algorithm on a real signal when the assumption of random Gaussian sensing vectors do not hold, we consider another example from [5]. An image of Stanford main quad of size 320×1280 pixels is used, and measurements are obtained via a set of L coded

Figure 3: Results for Example 3: Recovered image after (top) initialization (50 power iterations), (center) 1 pass using TWF, (bottom) 1 pass using ITWF

diffraction patterns as

$$\boldsymbol{y}^{(l)} = (\boldsymbol{F}\boldsymbol{D}^{(l)}\boldsymbol{x})^* \odot (\boldsymbol{F}\boldsymbol{D}^{(l)}\boldsymbol{x}), \quad 1 \le l \le L,$$

where $\boldsymbol{F}$ is the DFT matrix and $\boldsymbol{D}^{(l)}$ is a diagonal matrix (representing a mask) containing independent entries, each uniformly distributed over $\{+1, -1, j, -j\}$. The number of measurements is nL with $L = 12$, and the notation $|\cdot|^2$ denotes elementwise magnitude squared.

As in [5], we initialize the algorithm with 50 truncated power iterations, at the end of which the relative root mean squared error (rmse) is 0.35. Another pass through the data using TWF gets it down to 0.15. Increasing the number of passes to 2, 5 and 10 achieves 7.1×10^{-2}, 2.3×10^{-2} and 6.6×10^{-3} respectively. Making one pass through the data using ITWF gets it down from 0.35 to 8.2×10^{-3}. Increasing the number of passes to 2, 5 and 10 achieves 3.0×10^{-4}, 1.4×10^{-8} and 9.0×10^{-16} respectively.

Remark: The sensing matrices in this example are highly structured (DFT, diagonal), due to which computing the sum of the n gradients for one mask can be accomplished with $O(n \log n)$ computations instead of $O(n^2)$, by utilizing FFT algorithms. Hence, the motivation for ITWF that one iteration can be made m-times cheaper does not hold in this case. Hence we consider an *increment* to be the set of measurements corresponding to one mask (L measurements) instead of one measurement. This means that each iteration of ITWF is L-times cheaper than that of TWF where L is the total number of masks; thus one iteration of TWF is computationally equivalent to L iterations of ITWF.

4.4 Example 4: Noisy measurements

In this example, we generate noisy measurements as follows. As before, the sensing vectors are chosen to be i.i.d. $\mathcal{N}(0, I)$. The measurements are generated according to (11). Theorem 2 effectively says that ITWF achieves

$$\frac{\text{dist}^2(\boldsymbol{z}^{(T)}, \boldsymbol{x})}{\|\boldsymbol{x}\|^2} \sim \frac{1}{\text{SNR}}, \tag{13}$$

where SNR is the signal-to-noise ratio $\frac{\sum_{i=1}^m (\boldsymbol{a}_i^T \boldsymbol{x})^4}{\sum_{i=1}^m \eta_i^2}$, which is approximately $\frac{3m\|\boldsymbol{x}\|^4}{\|\boldsymbol{\eta}\|^2}$.

Under the Poisson model, since $\sum_{i=1}^m \eta_i^2 \approx \sum_{i=1}^m |\boldsymbol{a}_i^T \boldsymbol{x}|^2 \approx m\|\boldsymbol{x}\|^2$, we refer to $3\|\boldsymbol{x}\|^2$ as the SNR. Comparing the final relative MSE of TWF and ITWF at various values of SNR Figure 4 shows that the final relative MSE (LHS of (13)) of ITWF is in fact nearly equal to that of ITWF at all values of SNR.

The gain in computational complexity provided by ITWF that we observed in the noiseless case (ref. Figure 2) is however abated in the noisy case, as can be seen from the left panel in Figure 5. Experimenting with a diminishing step-size rule (step-size in ℓth pass chosen as $\Theta\left(\frac{1}{n\ell}\right)$), we find that it not only results in a small final MSE, but also reduces the number of passes it takes for ITWF to reach close to the final MSE by allowing us to exploit a larger step-size in the initial passes, as shown in the right panel in Figure 5.

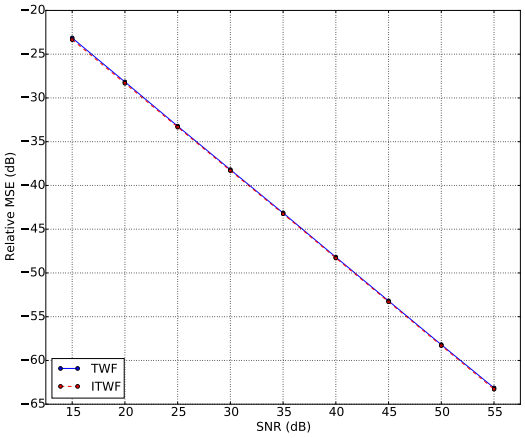

Figure 4: Final Relative MSE achieved as a function of SNR

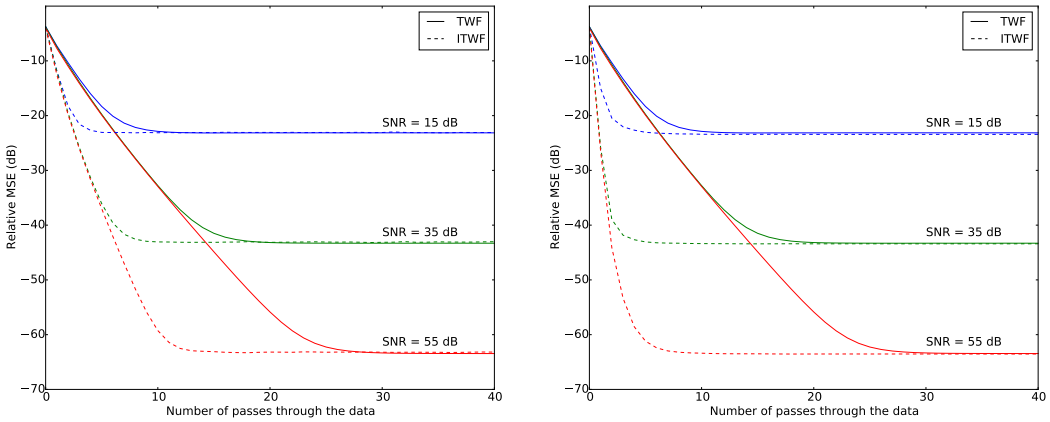

Figure 5: Relative MSE achieved by TWF (step size 0.2) compared with (left) ITWF with constant step-size and (right) ITWF with a decreasing step size

5 Remarks and Future Directions

While loosely bounding the constants arising during the analysis provides a recommendation for the step-size (for the random Gaussian model) which is $\approx \frac{0.00016}{n}$, this can be slightly pessimistic. In our numerical experiments, we find that larger step sizes also work.

Developing incremental methods that are able to extract the benefits of truncation for other problems such as sparse phase retrieval, low-rank matrix recovery from linear/quadratic measurements would be highly interesting.

References

[1] Emmanuel Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

[2] Emmanuel Candes and Xiaodong Li. Solving quadratic equations via phaselift when there are about as many equations as unknowns. *Foundations of Computational Mathematics*, 14(5):1017–1026, 2014.

[3] Emmanuel Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, April 2015.

[4] Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi. Phase retrieval using alternating minimization. *IEEE Transactions on Signal Processing*, 63(18):4814–4826, Sept 2015.

[5] Yuxin Chen and Emmanuel Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems 28*, pages 739–747. 2015.

[6] Kishore Jaganathan, Samet Oymak, and Babak Hassibi. Sparse phase retrieval: Uniqueness guarantees and recovery algorithms. *CoRR*, abs/1311.2745, 2013.

[7] Ramtin Pedarsani, Kangwook Lee, and Kannan Ramchandran. Phasecode: Fast and efficient compressive phase retrieval based on sparse-graph-codes. *CoRR*, abs/1408.0034, 2014.

[8] Emmanuel Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Applied and Computational Harmonic Analysis*, 39(2):277 – 299, 2015.

[9] Kishore Jaganathan, Yonina Eldar, and Babak Hassibi. STFT phase retrieval: Uniqueness guarantees and recovery algorithms. *CoRR*, abs/1508.02820, 2015.

[10] Kishore Jaganathan, Yonina Eldar, and Babak Hassibi. Phase retrieval: An overview of recent developments. *CoRR*, abs/1510.07713, 2015.

[11] Ohad Shamir. A stochastic PCA and SVD algorithm with an exponential convergence rate. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015*, pages 144–152, 2015.

[12] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, pages 315–323, 2013.

[13] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *CoRR*, abs/1011.3027, 2010.

6 Appendix

6.1 Proof of Theorem 1

The proof requires introducing some constants, and we collect their definitions before proceeding to the proof.

Define

$$\zeta_1 := \max \left\{ \begin{array}{c} \mathbb{E}\left[\xi^2 \mathbb{1}_{|\xi| \notin (\sqrt{1.01}\alpha_z^{\text{lb}}, \sqrt{0.99}\alpha_z^{\text{ub}})}\right], \\ \mathbb{E}\left[\mathbb{1}_{|\xi| \notin (\sqrt{1.01}\alpha_z^{\text{lb}}, \sqrt{0.99}\alpha_z^{\text{ub}})}\right] \end{array} \right\},$$

and

$$\zeta_2 := \max \left\{ \begin{array}{c} \mathbb{E}\left[\xi^2 \mathbb{1}_{|\xi| \geq 0.98\alpha_x}\right], \\ \mathbb{E}\left[\mathbb{1}_{|\xi| \geq 0.98\alpha_x}\right] \end{array} \right\},$$

and

$$\zeta_3 := \mathbb{E}\left[\xi^2 \mathbb{1}_{|\xi| \geq \sqrt{0.99}\gamma}\right],$$

where $\xi \sim \mathcal{N}(0,1)$ and γ, α_x, α_z^{ub} and α_z^{lb} can be chosen to be 5, 0.3, 5 and 5 respectively.

The term δ_{init} appearing in the proofs can be considered to be equal to 0.1 for concreteness.

The intermediate lemmas and propositions will establish statements that hold for all $\boldsymbol{z}$ that are in a neighborhood of $\boldsymbol{x}$ as follows:

$$\frac{\text{dist}(\boldsymbol{z}, \boldsymbol{x})}{\|\boldsymbol{z}\|} \leq \min \left\{ \delta_{\text{init}}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98(\alpha_z^{\text{lb}})^5}}{\sqrt{(\alpha_z^{\text{ub}} + 1.01\alpha_x(1+\delta_{\text{init}}))^3}} \right\}, \tag{14}$$

Note that the initialization stage can be used to guarantee that $\boldsymbol{z}^{(0)}$ is in this neighborhood of $\boldsymbol{x}$.

Then, to prove the theorem, it suffices to prove Proposition 1. The reason why proving this proposition suffices is as follows. Consider the statement (15) in Proposition 1. We can take expectation of the RHS with respect to i_{t-1}, and then apply Proposition 1 again to get a similar relation for the previous iteration. Continuing this process till we arrive at the initialization point $\boldsymbol{z}^{(0)}$, and noting that the initialization stage guarantees that $\text{dist}(\boldsymbol{z}^{(0)}, \boldsymbol{x})$ is at most a small constant times $\|\boldsymbol{x}\|$, we arrive at the statement of Theorem 1. This concludes the proof of Theorem 1. □

Proposition 1. *When $\{\boldsymbol{a}_i\}_{i=1}^m$ are independent $\mathcal{N}(0, I)$, there exist universal constants $C, c_0, c_1, c_2, \rho > 0$, such that with probability at least $1 - Cm\exp(-c_1 n)$ and $\mu = c_2/n$, we have*

$$\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)}, \boldsymbol{x})\right] \leq \left(1 - \frac{\rho}{n}\right) \cdot \text{dist}^2(\boldsymbol{z}^{(t)}, \boldsymbol{x}), \tag{15}$$

holds simultaneously for all $\boldsymbol{z}^{(t)} \in \mathbb{R}^n$ satisfying (14)*, if $m \geq c_0 n$.*

Proof. Let $\boldsymbol{h}$ denote $\boldsymbol{z}^{(t)} - \boldsymbol{x}$, and $\mathcal{E}_t^i = \mathcal{E}_{1,t}^i \cap \mathcal{E}_3^i$ denote the overall truncation. The ITWF update (7), after some simple algebraic manipulations, gives us

$$\begin{aligned} \mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)}, \boldsymbol{x})\right] = \|\boldsymbol{h}\|^2 &- \frac{4\mu}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)\mathbb{1}_{\mathcal{E}_t^i} \\ &+ \frac{4\mu^2}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2 |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)^2 \mathbb{1}_{\mathcal{E}_t^i} \end{aligned} \tag{16}$$

Lemmas 1, 2 and 3 show that for appropriate ζ_1, ζ_2 and ζ_3 and any small $\delta > 0$, the following two relations hold *simultaneously* for all $\boldsymbol{z}$, $\boldsymbol{h}$ and $\boldsymbol{x}$ with high probability:

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)\mathbb{1}_{\mathcal{E}_t^i} \geq (0.99 - \zeta_1 - \zeta_2 - \zeta_3 - 3\delta)\|\boldsymbol{h}\|^2,$$

and

$$\frac{1}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2 |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)^2 \mathbb{1}_{\mathcal{E}_t^i} \leq c_4 n(1+\delta)\|\boldsymbol{h}\|^2,$$

if m is $O(n)$ and c_4 is a constant. The proofs mainly involve the application of concentration bounds and ϵ-net arguments. Substituting these relations in (16) and choosing μ to be say $\frac{0.99-\zeta_1-\zeta_2-\zeta_3-3\delta}{2c_4(1+\delta)n}$, we get that

$$\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)},\boldsymbol{x})\right] \leq \left(1-\frac{\rho}{n}\right)\text{dist}^2(\boldsymbol{z}^{(t)},\boldsymbol{x}).$$

□

Lemma 1. *For any $\delta > 0$, there exist universal constants $C, c_0, c_1 > 0$ such that if $m > c_1 n\delta^{-2}\log(1/\delta)$, then*

$$\frac{1}{m}\sum_{i=1}^{m}|\boldsymbol{a}_i^T\boldsymbol{h}|^2\mathbb{1}_{\mathcal{E}_{1,t}^i}\mathbb{1}_{\mathcal{E}_3^i} \geq (1-\zeta_1-\zeta_2-\zeta_3-2\delta)\|\boldsymbol{h}\|^2,$$

with probability $1-C\exp(-c_0 m\delta^2)$, simultaneously for all non-zero vectors $\boldsymbol{h},\boldsymbol{z},\boldsymbol{x}\in\mathbb{R}^n$.

Proof. The proof follows a similar path as Lemma 4 in [5]. We can restrict attention to unit norm $\boldsymbol{h}$, $\boldsymbol{z}$ and $\boldsymbol{x}$. We first assume that $\boldsymbol{z}$, $\boldsymbol{h}$ and $\boldsymbol{x}$ are independent from $\{\boldsymbol{a}_i\}_i$. Later, by using an ϵ-net argument, we extend it to all unit norm vectors $\boldsymbol{z}$, $\boldsymbol{h}$ and $\boldsymbol{x}$.

First, we note that with high probability, we have

$$\begin{aligned}\mathbb{1}_{\mathcal{E}_3^i} = \mathbb{1}_{\left\{\frac{y_i}{\frac{1}{m}\sum_{i=1}^m y_i}\leq\alpha_x^2\right\}} &\geq \mathbb{1}_{\left\{\frac{|\boldsymbol{a}_i^T\boldsymbol{x}|}{\|\boldsymbol{x}\|}\leq\alpha_x(1-0.01)\right\}}\\ &= \mathbb{1}_{\left\{\frac{|\boldsymbol{a}_i^T\boldsymbol{x}|}{\|\boldsymbol{x}\|}\leq 0.99\alpha_x\right\}},\end{aligned}$$

where the inequality follows from Theorem 5.39 in [13], which says that

$$\left\|\frac{1}{m}\sum_{i=1}^{m}\boldsymbol{a}_i\boldsymbol{a}_i^T - I\right\| \leq \delta, \tag{17}$$

with probability $1-C\exp(-c_1 m\delta^2)$ if $m\geq c_0 n\delta^{-2}$. For convenience, let us denote $0.99\alpha_x$ by $\widetilde{\alpha}_x$. In the remainder of this proof and the rest of the proofs for the noiseless case, we will use the established lower bound

$$\mathbb{1}_{\mathcal{E}_3^i} \geq \mathbb{1}_{\left\{\frac{|\boldsymbol{a}_i^T\boldsymbol{x}|}{\|\boldsymbol{x}\|}\leq\widetilde{\alpha}_x\right\}}$$

implicitly whenever required.

Now define the functions

$$\chi_z(\tau) = \begin{cases}1, & \text{if } |\tau|\in[\sqrt{1.01}\alpha_z^{\text{lb}},\sqrt{0.99}\alpha_z^{\text{ub}}]\\ 100\left(1-\left(\frac{\tau}{\alpha_z^{\text{ub}}}\right)^2\right), & \text{if } |\tau|\in[\sqrt{0.99}\alpha_z^{\text{ub}},\alpha_z^{\text{ub}}]\\ 100\left(\left(\frac{\tau}{\alpha_z^{\text{lb}}}\right)^2-1\right), & \text{if } |\tau|\in[\alpha_z^{\text{lb}},\sqrt{1.01}\alpha_z^{\text{lb}}]\\ 0, & \text{else,}\end{cases}$$

and

$$\chi_x(\tau) = \begin{cases}1, & \text{if } |\tau|\leq\sqrt{0.99}\widetilde{\alpha}_x\\ 100\left(1-\left(\frac{\tau}{\widetilde{\alpha}_x}\right)^2\right), & \text{if } |\tau|\in[\sqrt{0.99}\widetilde{\alpha}_x,\widetilde{\alpha}_x]\\ 0, & \text{else,}\end{cases}$$

and

$$\chi_h(\tau) = \begin{cases}1, & \text{if } |\tau|\leq\sqrt{0.99}\gamma\\ 100\left(1-\left(\frac{\tau}{\gamma}\right)^2\right), & \text{if } |\tau|\in[\sqrt{0.99}\gamma,\gamma]\\ 0, & \text{else.}\end{cases}$$

Since $\boldsymbol{h}$, $\boldsymbol{x}$ and $\boldsymbol{z}$ are assumed to be unit vectors, we have

$$\mathbb{1}_{\mathcal{E}^i_{1,t}}\mathbb{1}_{\mathcal{E}^i_3} \geq \mathbb{1}_{\mathcal{E}^i_{1,t}}\mathbb{1}_{\mathcal{E}^i_3}\mathbb{1}_{\mathcal{E}^i_{4,t}} \geq \chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}) \geq 0,$$

where $\mathcal{E}^i_{4,t}$ is defined to be the event $\{|\boldsymbol{a}_i^T\boldsymbol{h}| \leq \gamma\|\boldsymbol{h}\|\}$, and γ can be chosen to be 5. This means that

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\mathbb{1}_{\mathcal{E}^i_{1,t}}\mathbb{1}_{\mathcal{E}^i_3} \geq \frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}). \tag{18}$$

The reason for introducing these functions is that their $O(1)$-Lipschitz continuity allows us to apply the ϵ-net argument later in the proof. The event $\mathcal{E}^i_{4,t}$ allows us to discard atypically large terms in the summation, which is useful while applying the ϵ-net argument.

We now lower bound the RHS of the above inequality. The expectation of each term in the summation can be lower bounded as follows:

$$\begin{aligned}
&\mathbb{E}\left[|\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h})\right] \\
&\quad\geq \mathbb{E}\left[|\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\right] + \mathbb{E}\left[|\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\right] + \mathbb{E}\left[|\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_h(\boldsymbol{a}_i^T\boldsymbol{h})\right] - 2\mathbb{E}\left[|\boldsymbol{a}_i^T\boldsymbol{h}|^2\right] \\
&\quad\geq (1-\zeta_1)\|\boldsymbol{h}\|^2 + (1-\zeta_2)\|\boldsymbol{h}\|^2 + (1-\zeta_3)\|\boldsymbol{h}\|^2 - 2\|\boldsymbol{h}\|^2 \\
&\quad= (1-\zeta_1-\zeta_2-\zeta_3)\|\boldsymbol{h}\|^2,
\end{aligned}$$

where these inequalities follow by similar arguments as equation (120) in [5].

Observing that $|\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h})$ is a sub-exponential random variable with sub-exponential norm $O(\|\boldsymbol{h}\|^2)$, we can use Proposition 5.16 from [13] to get that with probability $1-\exp(-\Omega(\delta^2 m))$,

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}) \geq (1-\zeta_1-\zeta_2-\zeta_3-\delta)\|\boldsymbol{h}\|^2.$$

We now construct an ϵ-net $\mathcal{N}_\epsilon$, such that for any $(\boldsymbol{h},\boldsymbol{z},\boldsymbol{x})$ with $\|\boldsymbol{h}\| = \|\boldsymbol{z}\| = \|\boldsymbol{x}\| = 1$, there exists $(\boldsymbol{h}_0,\boldsymbol{z}_0,\boldsymbol{x}_0) \in \mathcal{N}_\epsilon$ such that $\|\boldsymbol{h}-\boldsymbol{h}_0\| \leq \epsilon$, $\|\boldsymbol{z}-\boldsymbol{z}_0\| \leq \epsilon$ and $\|\boldsymbol{x}-\boldsymbol{x}_0\| \leq \epsilon$, and $|\mathcal{N}_\epsilon| \leq \left(1+\frac{2}{\epsilon}\right)^{3n}$. Taking a union bound over this set gives us that

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}_0) \geq (1-\zeta_1-\zeta_2-\zeta_3-\delta)\|\boldsymbol{h}_0\|^2,$$

holds for all $(\boldsymbol{h}_0,\boldsymbol{z}_0,\boldsymbol{x}_0) \in \mathcal{N}_\epsilon$ with probability a least $1-\left(1+\frac{2}{\epsilon}\right)^{3n}\exp(-\Omega(\delta^2 m))$.

Now, since $\chi_z(\sqrt{\tau})$, $\chi_x(\sqrt{\tau})$ and $\tau\chi_h(\sqrt{\tau})$ are Lipschitz functions with Lipschitz constant $O(1)$, we get that for any unit vector pair $\boldsymbol{h}$, $\boldsymbol{z}$ and $\boldsymbol{x}$,

$$\begin{aligned}
&\left||\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x})\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}) - |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}_0)\right| \\
&\quad\leq |\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)|\cdot\left||\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}) - |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}_0)\right| \\
&\qquad + |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2|\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}_0)\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)|\cdot\left|\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}) - \chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)\right| \\
&\qquad + |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\chi_h(\boldsymbol{a}_i^T\boldsymbol{h}_0)|\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)|\cdot\left|\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}) - \chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)\right| \\
&\quad\lesssim \left||\boldsymbol{a}_i^T\boldsymbol{h}|^2 - |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\right| + \left||\boldsymbol{a}_i^T\boldsymbol{z}|^2 - |\boldsymbol{a}_i^T\boldsymbol{z}_0|^2\right| + \left||\boldsymbol{a}_i^T\boldsymbol{x}|^2 - |\boldsymbol{a}_i^T\boldsymbol{x}_0|^2\right|.
\end{aligned}$$

Hence, there exists a universal constant c_3 such that

$$\begin{aligned}
&\left|\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z})\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}) - \frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\chi_z(\boldsymbol{a}_i^T\boldsymbol{z}_0)\chi_x(\boldsymbol{a}_i^T\boldsymbol{x}_0)\right| \\
&\quad\lesssim \frac{1}{m}\sum_{i=1}^m \left(\left||\boldsymbol{a}_i^T\boldsymbol{h}|^2 - |\boldsymbol{a}_i^T\boldsymbol{h}_0|^2\right| + \left||\boldsymbol{a}_i^T\boldsymbol{z}|^2 - |\boldsymbol{a}_i^T\boldsymbol{z}_0|^2\right| + \left||\boldsymbol{a}_i^T\boldsymbol{x}|^2 - |\boldsymbol{a}_i^T\boldsymbol{x}_0|^2\right|\right) \\
&\quad\leq 3c_3\epsilon,
\end{aligned}$$

where the last inequality follows by invoking Lemma 1 and Lemma 2 from [5], assuming $\epsilon < 1/2$. By choosing $\epsilon = \frac{\delta}{3c_3}$ and recalling (18), we get the lemma. □

Lemma 2. *Consider any $\boldsymbol{z} \in \mathbb{R}^n$, and $\boldsymbol{h} := \boldsymbol{z} - \boldsymbol{x}$ such that*

$$\frac{\|\boldsymbol{h}\|}{\|\boldsymbol{z}\|} \leq \min\left\{\delta_{\text{init}}, \frac{\alpha_z^{\text{lb}}}{6}, \frac{\sqrt{98(\alpha_z^{\text{lb}})^5}}{\sqrt{(\alpha_z^{ub} + 1.01\alpha_x(1+\delta_{init}))^3}}\right\}.$$

Then, for any $\delta \in (0,1)$, there exist universal constants $C, c_0, c_1 > 0$ such that if $m \geq c_1 n\delta^{-2}\log(1/\delta)$, it holds that

$$\frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}} \mathbb{1}_{\mathcal{E}_t^i} \geq -(0.01+\delta)\|\boldsymbol{h}\|^2,$$

with probability at least $1 - C\exp(-c_0\delta^2 m)$.

Proof. We first note that

$$\begin{aligned}
&\frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}} \mathbb{1}_{\mathcal{E}_t^i} \\
&\qquad \geq \frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}} \mathbb{1}_{\mathcal{E}_t^i \cap \mathcal{D}_t^i} \\
&\qquad = -\left(\frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left|\frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right| \mathbb{1}_{\mathcal{E}_t^i \cap \mathcal{D}_t^i}\right),
\end{aligned}$$

where

$$\mathcal{D}_t^i = \left\{\frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}} \leq 0\right\}.$$

Also note that $\mathcal{D}_t^i \Leftrightarrow \widetilde{\mathcal{D}}_t^i$, where

$$\widetilde{\mathcal{D}}_t^i = \left\{|\boldsymbol{a}_i^T\boldsymbol{h}| \geq |\boldsymbol{a}_i^T\boldsymbol{z}|\right\} \cap \left\{|\boldsymbol{a}_i^T\boldsymbol{h}| \geq |\boldsymbol{a}_i^T\boldsymbol{x}|\right\}.$$

So, we now proceed to obtain an upper bound on $\frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left|\frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right| \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i}$. We have

$$\begin{aligned}
&\frac{1}{m}\sum_{i=1}^{m} |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left|\frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right| \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \\
&\qquad \leq \frac{1}{m}\sum_{i=1}^{m} \frac{|\boldsymbol{a}_i^T\boldsymbol{h}|^3}{|\boldsymbol{a}_i^T\boldsymbol{z}|} \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \\
&\qquad \leq \frac{1}{m}\sum_{i=1}^{m} \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \max_{i:\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \left\{\frac{|\boldsymbol{a}_i^T\boldsymbol{h}|^3}{|\boldsymbol{a}_i^T\boldsymbol{z}|}\right\}.
\end{aligned}$$

Now, we can use Lemma 6 from [5] to upper bound $\sum_{i=1}^{m} \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i}$ with high probability, as long as m is sufficiently large. Assuming that γ is a real number that is at least 2 and at most $\frac{\alpha_z^{\text{lb}}\|\boldsymbol{z}\|}{\|\boldsymbol{h}\|}$, this is done as follows:

$$\begin{aligned}
\frac{1}{m}\sum_{i=1}^{m} \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} &\leq \frac{1}{m}\sum_{i=1}^{m} \mathbb{1}_{\{|\boldsymbol{a}_i^T\boldsymbol{h}| \geq \gamma\|\boldsymbol{h}\|\}} \\
&\leq \frac{1}{0.49\gamma}\exp(-0.485\gamma^2) + \frac{\epsilon}{\gamma^2} \\
&\leq \frac{1}{9800}\left(\frac{\|\boldsymbol{h}\|}{\alpha_z^{\text{lb}}\|\boldsymbol{z}\|}\right)^4 + \frac{\epsilon}{(\alpha_z^{\text{lb}})^2}\left(\frac{\|\boldsymbol{h}\|}{\|\boldsymbol{z}\|}\right)^2,
\end{aligned}$$

which holds simultaneously for all $\boldsymbol{h}$ with probability $1 - C\exp(-c_0\epsilon^2 m)$ if $m \geq c_1 n\epsilon^{-2}\log(1/\epsilon)$, where the last step also requires $\frac{\|\boldsymbol{h}\|}{\|\boldsymbol{z}\|} \leq \frac{\alpha_z^{\text{lb}}}{6}$.

Since $|\boldsymbol{a}_i^T\boldsymbol{h}| = |\boldsymbol{a}_i^T\boldsymbol{z} - \boldsymbol{a}_i^T\boldsymbol{x}| \le |\boldsymbol{a}_i^T\boldsymbol{z}| + |\boldsymbol{a}_i^T\boldsymbol{x}|$ and by the definition of $\mathcal{E}_t^i$, we have

$$\max_{i:\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \left\{ \frac{|\boldsymbol{a}_i^T\boldsymbol{h}|^3}{|\boldsymbol{a}_i^T\boldsymbol{z}|} \right\} \le \frac{(\alpha_z^{\text{ub}} + 1.01\alpha_x(1+\delta_{\text{init}}))^3}{\alpha_z^{\text{lb}}} \|\boldsymbol{z}\|^2$$

if we assume $\|\boldsymbol{x}\| \le (1+\delta_{\text{init}})\|\boldsymbol{z}\|$, which is true if $\frac{\|\boldsymbol{h}\|}{\|\boldsymbol{z}\|} \le \delta_{\text{init}}$.

As a result of these two upper bounds, we get

$$\begin{aligned}
&\frac{1}{m}\sum_{i=1}^m \mathbb{1}_{\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \max_{i:\mathcal{E}_t^i \cap \widetilde{\mathcal{D}}_t^i} \left\{ \frac{|\boldsymbol{a}_i^T\boldsymbol{h}|^3}{|\boldsymbol{a}_i^T\boldsymbol{z}|} \right\} \\
&\quad \le \left\{ \frac{(\alpha_z^{\text{ub}} + 1.01\alpha_x(1+\delta_{\text{init}}))^3}{9800(\alpha_z^{\text{lb}})^5} \frac{\|\boldsymbol{h}\|^2}{\|\boldsymbol{z}\|^2} + \frac{\epsilon(\alpha_z^{\text{ub}} + 1.01\alpha_x(1+\delta_{\text{init}}))^3}{(\alpha_z^{\text{lb}})^3} \right\} \|\boldsymbol{h}\|^2.
\end{aligned}$$

Now, by choosing ϵ appropriately, and if

$$\frac{\|\boldsymbol{h}\|}{\|\boldsymbol{z}\|} \le \frac{\sqrt{98(\alpha_z^{\text{lb}})^5}}{\sqrt{(\alpha_z^{\text{ub}} + 1.01\alpha_x(1+\delta_{\text{init}}))^3}},$$

we get the claim of the lemma. □

Lemma 3. *For all non-zero vectors $\boldsymbol{z} \in \mathbb{R}^n$ and $\boldsymbol{h} := \boldsymbol{z} - \boldsymbol{x}$ with $\|\boldsymbol{h}\| \le \delta_{\text{init}}\|\boldsymbol{z}\|$, we have if $m \ge c_1 n\delta^{-2}$,*

$$\frac{1}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2 |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)^2 \mathbb{1}_{\mathcal{E}_t^i} \le c_4(1+\delta)n\|\boldsymbol{h}\|^2,$$

with probability $1 - Cm\exp(-c_0 n\delta^2)$.

Proof. We have that $\|\boldsymbol{a}_i\| \le \sqrt{6n}$ for all i with probability $1 - m\exp(-1.5n)$. The following chain of inequalities uses this, and the definition of $\mathcal{E}_t^i$.

$$\begin{aligned}
&\frac{1}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2 |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)^2 \mathbb{1}_{\mathcal{E}_t^i} \\
&\le \frac{1}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2 |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \left(1 + \frac{1.01\alpha_x(1+\delta_{\text{init}})}{\alpha_z^{\text{lb}}}\right)^2 \mathbb{1}_{\mathcal{E}_t^i} \\
&\le 6n\left(1 + 1.01\frac{\alpha_x(1+\delta_{\text{init}})}{\alpha_z^{\text{lb}}}\right)^2 \frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2 \\
&\le c_4 n(1+\delta)\|\boldsymbol{h}\|^2,
\end{aligned}$$

where the last inequality follows since by Theorem 5.39 in [13],

$$\left\| \frac{1}{m}\sum_{i=1}^m \boldsymbol{a}_i\boldsymbol{a}_i^T - I \right\| \le \delta,$$

with probability $1 - C\exp(-c_1 m\delta^2)$ if $m \ge c_0 n\delta^{-2}$. We have also used the fact that $y_i \le \alpha_x^2(\frac{1}{m}\sum_{i=1}^m y_i)$ implies

$$|\boldsymbol{a}_i^T\boldsymbol{x}| \le 1.01\alpha_x\|\boldsymbol{x}\| \le 1.01\alpha_x(1+\delta_{\text{init}})\|\boldsymbol{z}\|.$$

It can be seen that c_4 stands for the quantity $6\left(1 + \frac{1.01\alpha_x(1+\delta_{\text{init}})}{\alpha_z^{\text{lb}}}\right)^2$. □

6.2 Proof of Theorem 2

In the noisy case, the measurements are generated according to $y_i = |\boldsymbol{a}_i^T \boldsymbol{x}|^2 + \eta_i$. Recall that we are using $\mathcal{E}_t^i$ to denote $\mathcal{E}_{1,t}^i \cap \mathcal{E}_3^i$.

Proceeding in a similar manner as Proposition 1, and recalling that we are using $\boldsymbol{h} = \boldsymbol{z}^{(t)} - \boldsymbol{x}$, we get that

$$\begin{aligned}
&\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)}, \boldsymbol{x})\right] \\
&= \|\boldsymbol{h}\|^2 - \frac{4\mu}{m}\sum_{i=1}^m \boldsymbol{a}_i^T\boldsymbol{h}\left(\frac{|\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}|^2 - |\boldsymbol{a}_i^T\boldsymbol{x}|^2 - \eta_i}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)\mathbb{1}_{\mathcal{E}_t^i} \\
&\quad + \frac{4\mu^2}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2\left(\frac{|\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}|^2 - |\boldsymbol{a}_i^T\boldsymbol{x}|^2 - \eta_i}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)^2\mathbb{1}_{\mathcal{E}_t^i} \\
&= \|\boldsymbol{h}\|^2 - \frac{4\mu}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)\mathbb{1}_{\mathcal{E}_t^i} + \frac{4\mu}{m}\sum_{i=1}^m \frac{(\boldsymbol{a}_i^T\boldsymbol{h})\eta_i}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\mathbb{1}_{\mathcal{E}_t^i} \\
&\quad + \frac{4\mu^2}{m}\sum_{i=1}^m \|\boldsymbol{a}_i\|^2|\boldsymbol{a}_i^T\boldsymbol{h}|^2\left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)^2\mathbb{1}_{\mathcal{E}_t^i} + \frac{4\mu^2}{m}\sum_{i=1}^m \frac{\eta_i^2\|\boldsymbol{a}_i\|^2}{|\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}|^2}\mathbb{1}_{\mathcal{E}_t^i} \\
&\quad - \frac{8\mu^2}{m}\sum_{i=1}^m \frac{(\boldsymbol{a}_i^T\boldsymbol{h})\eta_i}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\left(1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right)\|\boldsymbol{a}_i\|^2\mathbb{1}_{\mathcal{E}_t^i} \qquad (19)
\end{aligned}$$

To shorten the expressions, we will resort to denoting $\boldsymbol{z}^{(t)}$ by $\boldsymbol{z}$ in the remainder of this proof. The follow lemma is the noisy-case counterpart of Lemma 1.

Lemma 4. *For any $\boldsymbol{h}$, $\boldsymbol{z}$ and $\boldsymbol{x}$, for any $\epsilon, \delta > 0$, there exist constants $C, c_0, c_1 > 0$ such that if $m > c_1 n\delta^{-2}\log(1/\delta)$, then*

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\mathbb{1}_{\mathcal{E}_t^i} \geq (1 - \zeta_1 - \zeta_2 - \zeta_3 - 2\delta)\|\boldsymbol{h}\|^2,$$

with probability $1 - C\exp(-c_0 m\delta^2)$, for appropriately redefined $\zeta_1, \zeta_2, \zeta_3$.

Proof. We have for any i,

$$\begin{aligned}
\mathbb{1}_{\mathcal{E}_3^i} = \mathbb{1}_{\left\{y_i \leq \alpha_x^2\left(\frac{1}{m}\sum_{i=1}^m y_i\right)\right\}} &\geq \mathbb{1}_{\left\{|\boldsymbol{a}_i^T\boldsymbol{x}|^2 + |\eta_i| \leq \alpha_x^2\left((1-0.01)\|\boldsymbol{x}\|^2 - \frac{1}{m}\|\boldsymbol{\eta}\|_1\right)\right\}} \\
&\geq \mathbb{1}_{\left\{|\boldsymbol{a}_i^T\boldsymbol{x}|^2 + \|\boldsymbol{\eta}\|_\infty \leq \alpha_x^2(0.99\|\boldsymbol{x}\|^2 - \|\boldsymbol{\eta}\|_\infty)\right\}} \\
&\geq \mathbb{1}_{\left\{|\boldsymbol{a}_i^T\boldsymbol{x}|^2 + \epsilon_\eta\|\boldsymbol{x}\|^2 \leq \alpha_x^2(0.99\|\boldsymbol{x}\|^2 - \epsilon_\eta\|\boldsymbol{x}\|^2)\right\}} \\
&= \mathbb{1}_{\left\{|\boldsymbol{a}_i^T\boldsymbol{x}|^2 \leq (\alpha_x^2(0.99 - \epsilon_\eta) - \epsilon_\eta)\|\boldsymbol{x}\|^2\right\}}
\end{aligned}$$

where the inequalities follow by (17) and the assumption $\|\boldsymbol{\eta}\|_\infty \leq \epsilon_\eta\|\boldsymbol{x}\|^2$. Thus, we get that the RHS of the above expression is similar to $\mathbb{1}_{\left\{\frac{|\boldsymbol{a}_i^T\boldsymbol{x}|}{\|\boldsymbol{x}\|} \leq \widetilde{\alpha}_x\right\}}$. Then, by applying similar arguments as in Lemma 1, with appropriately redefined ζ_1, ζ_2 and ζ_3, we can conclude that the following holds with high probability,

$$\frac{1}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\mathbb{1}_{\mathcal{E}_t^i} \geq (1 - \zeta_1 - \zeta_2 - \zeta_3 - 2\delta)\|\boldsymbol{h}\|^2.$$

□

The following chain of inequalities, continuing from (19), uses the fact that $\|\boldsymbol{a}_i\|^2 \leq 6n$ holds with high probability, and

$$\left|1 + \frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}^{(t)}}\right| \leq \left(1 + \frac{\sqrt{\alpha_x^2(1.01 + \epsilon_\eta) + \epsilon_\eta}(1 + \delta_{\text{init}})}{\alpha_z^{\text{lb}}}\right) = \widetilde{c}_4,$$

which also holds with high probability by the definition of $\mathcal{E}_t^i$, by $\|\boldsymbol{h}\| \le \delta_{\text{init}}\|\boldsymbol{z}\|$ and by the assumption on the noise $\|\boldsymbol{\eta}\|_\infty \le \epsilon_\eta \|\boldsymbol{x}\|^2$. The Cauchy-Schwarz inequality is applied in terms involving $(\boldsymbol{a}_i^T\boldsymbol{h})\eta_i$.

$$\begin{aligned}
&\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)},\boldsymbol{x})\right] \\
&\le \|\boldsymbol{h}\|^2 - 4\mu(1-\zeta_1-\zeta_2-\zeta_3-2\delta)\|\boldsymbol{h}\|^2 - \frac{4\mu}{m}\sum_{i=1}^m |\boldsymbol{a}_i^T\boldsymbol{h}|^2\left(\frac{\boldsymbol{a}_i^T\boldsymbol{x}}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)\mathbb{1}_{\mathcal{E}_t^i} \\
&\quad + \frac{4\mu}{\alpha_z^{\text{lb}}}\frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}\|\boldsymbol{h}\| + 4\mu^2 n c_4(1+\delta)\|\boldsymbol{h}\|^2 + \frac{24\mu^2 n}{(\alpha_z^{\text{lb}})^2}\frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{z}\|^2} \\
&\quad + \frac{48\mu^2 n\widetilde{c}_4\sqrt{1+\delta}}{\alpha_z^{\text{lb}}}\|\boldsymbol{h}\|\frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}
\end{aligned} \tag{20}$$

The third term in the RHS can be shown to be no more than a small constant times $\|\boldsymbol{h}\|^2$ in a similar manner as in Lemma 2. Thus, we get by collecting all constants for brevity,

$$\begin{aligned}
\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)},\boldsymbol{x})\right] &\le \|\boldsymbol{h}\|^2 - 4\mu C\|\boldsymbol{h}\|^2 + \mu C'\frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}\|\boldsymbol{h}\| + \mu^2 nC''\|\boldsymbol{h}\|^2 \\
&\quad + \mu^2 nC'''\frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{z}\|^2} + \mu^2 nC''''\|\boldsymbol{h}\|\frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}.
\end{aligned} \tag{21}$$

6.2.1 Regime I

Regime I is defined to be the condition $\frac{c_3\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|} \le \|\boldsymbol{h}\| \le c_4\|\boldsymbol{x}\|$ where c_3 is some large constant. In this regime, it is possible to show that the distance to the solution decreases geometrically, as in the noiseless case.

Because $\frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|} \le \frac{1}{c_3}\|\boldsymbol{h}\|$ in regime 1, we can substitute this in (21) and choose an appropriate $\mu = \Theta\left(\frac{1}{n}\right)$, assuming c_3 is large enough, to arrive at

$$\mathbb{E}_{i_t}\left[\text{dist}^2(\boldsymbol{z}^{(t+1)},\boldsymbol{x})\right] \le \left(1-\frac{\rho}{n}\right)\|\boldsymbol{h}\|^2.$$

Thus, we have linear convergence in Regime 1.

6.2.2 Regime II

If $\boldsymbol{h}$ is not in Regime I, we have $\|\boldsymbol{h}\| \le \frac{c_3\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}$. While the distance to the solution is not guaranteed to decrease in this regime, it can be shown that by taking one step of ITWF, the expected next iterate either stays in Regime II or goes to Regime I, but not beyond, thus resulting in a bounded error, as given by Theorem 2. We employ similar arguments as in [5].

The expected norm of the step is

$$\begin{aligned}
\mathbb{E}_{i_t}\left[\left\|\mu\nabla\ell(y_{i_t}, |\boldsymbol{a}_{i_t}^T\boldsymbol{z}^{(t)}|^2)\cdot\mathbb{1}_{\mathcal{E}_t^{i_t}}\right\|\right] &= \mathbb{E}_{i_t}\left[\left\|\mu\left(\frac{|\boldsymbol{a}_i^T\boldsymbol{z}|^2 - |\boldsymbol{a}_i^T\boldsymbol{x}|^2 - \eta_i}{\boldsymbol{a}_i^T\boldsymbol{z}}\right)\boldsymbol{a}_i\cdot\mathbb{1}_{\mathcal{E}_t^{i_t}}\right\|\right] \\
&\lesssim \mu\sqrt{n}\left(\|\boldsymbol{h}\| + \frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}\right) \\
&\lesssim \frac{1}{\sqrt{n}}\left(\|\boldsymbol{h}\| + \frac{\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}\right).
\end{aligned}$$

If $\|\boldsymbol{h}\| \le \frac{c_3\|\boldsymbol{\eta}\|}{\sqrt{m}\|\boldsymbol{z}\|}$, we have that

$$\begin{aligned}
\mathbb{E}_{i_t}[\text{dist}^2(\boldsymbol{z}^{(t+1)},\boldsymbol{x})] &\lesssim \|\boldsymbol{h}\|^2 + \frac{1}{n}\left(\|\boldsymbol{h}\|^2 + \frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{z}\|^2}\right) \\
&\le C\frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{z}\|^2}.
\end{aligned}$$

Thus, if C is small, we stay in Regime II, else we have that

$$C\frac{\|\boldsymbol{\eta}\|^2}{m\|\boldsymbol{z}\|^2} \leq C\frac{\|\boldsymbol{\eta}\|_\infty^2}{\|\boldsymbol{z}\|^2} \leq \widetilde{C}\|\boldsymbol{x}\|^2,$$

thus ending up in Regime I, assuming $\frac{\|\boldsymbol{\eta}\|_\infty}{\|\boldsymbol{x}\|^2} = \epsilon_\eta$ is sufficiently small.

This concludes the proof of Theorem 2. □